Review Panel Summary - SNS Controls Final Design Review

The SNS Controls Final Design Review was held at LANL on January 16, 2001. The review panel consisted of Curt Hovater, Paul Corredoura, and Chris Ziomek. This document summarizes the individual comments from these three reviewers.

Specific Comments from Curt Hovater

Overall I believe that this is a robust design that will work. I think that most of the potential gotchas in the design have been addressed. Most of the issues from last summer have been answered. I think it is great that you are getting your documentation together and working from an actual requirements document. I still think that some of the numbers for latency/delay through the system are of concern, but the system is agile enough to overcome this if there is a problem. The LLRF design team is excellent and appears to have a good working relationship with one another. The working relationship with the customer Labs is good but could be improved.

Schedule

> This still appears aggressive I would not be surprised if it slips 6 months. Case in point the CSTM1-2 board is going to be 10-12 layers, with RF, Analog and digital. The first one will go in the garbage can. The second one will have cross talk problems and might end up in the garbage can. The third one may be just right! Add any of this to the DSP board, or motherboard also and there is your six months.

Cost

> It's high. Beside myself many folks from other labs and industry agree that it is high. Especially since this does not include the racks, VXI crates and Power pc.

Mark Prokop

> Mark seems to be everywhere in the design. He is the glue that holds this thing together. Is he still working with the high power folks too? He is going to be under intense pressure when the boards come into assist in the debug and troubleshooting of how many boards? What happens if he leaves? Is there some one assisting him with all of the VHDL?

Off Tangent

> I am little worried that flippant hallway comments can find there way into the LLRF design in places it does not belong. The example given is the phase monitor on the HPRF board. This board should be exclusively for machine protection. I believe any analog accelerator diagnostics should be incorporated into the RF/analog daughter board. In addition what may end up happening is that the commissioning folks wind up using the wrong tools (it becomes a crutch) when they should be focusing on the labor saving intelligent diagnostics that digital RF gives you, 'nuff said!

Mother Board

> Powers supplies: two are switchers. What are there switching speeds? Could this foul up any of the digital stuff on the other daughter cards?

RF/Analog/CPLD Board

Isolation…Isolation…. Isolation….Paul brings up a good point. Its okay to have RF components on both sides especially diagnostics. Make life easy on the EECAD designer! He will have his hands full enough.

HPRF Board

Most of my concerns would address SRF interlocks and I will make sure Dave is informed for the SC LLRF design review.

DSP Board

I am not a DSP engineer. Why don't you use a floating point DSP for this project? From what I know it can make life much easier when coding especially for custom low run designs. It also gives much more leverage if any problems arise with the firmware. Considering that you have covered all of the bases very well why skimp here?

Clock Module

Measure the power levels of all three VCO's & LO and mux back into ADC. I think this will be very useful for trouble shooting during operations. We have found this very useful when trouble shooting our MO system.

Board Stack Up

I would have an outside person (who knows what they are doing, not me!) look over your board stack up. How you handle grounds, power, signals, and bus lines will have a great impact on how many versions that you have to produce/throw away. Do you have a set of design rules that you give the EECAD designer?

Specific Comments from Paul Corredoura

1. I spent some time looking over the material from the review and initially thought about the limitations caused by the relatively long delay in the FIR filters. First of all the passband of the filters is 1.2 MHz. There was talk about an unwanted cavity mode at ~1MHz so as presently specified the FIR filters do not reject this mode at all. I played a bit with Matlab to see if I could make an IIR filter with comparable delay but a lower and sharper cutoff. Attached is a crude Matlab script comparing the existing FIR filter to a 4th order Cauer filter with the first deep notch placed at 1 MHz. Turns out the delay increases by 20 degrees at 200 kHz but it does filter better at 1 MHz. Then is began thinking why are the FIR filters there at all? They are not (I believe) functioning as part of the controller but simply as filtering after the ADCs. The controllers must have their own bandwidth limiting functions (to limit the control bandwidth to the stated 200 kHz right?). If these filters are there because of habit you should review if they are really necessary. You might find that you can remove them and increase your controller bandwidth significantly because of the decrease in loop delay. If the FIR is removed than a smaller gate array may work, saving cost. I also modeled the typical 2 sample averager to kill the fSample/2 tone caused by the I/Q sampling and any mixer offsets. If you choose to look at removing the long FIR make sure Sung-il does another full simulation with a resonance corresponding to the 1

MHz cavity mode. The loop response must have sufficient attenuation at this mode to prevent oscillations or proper delay added to insure this mode has negative feedback. By the way if you decide to decimate the long FIR output to save processing check out the INTFILT Matlab function. It shows how to make a decimating FIR filter with less resources by changing the tap weights every clock (easy to do in a DSP not sure about a gate array implementation). I went to a seminar and one speaker - Fred Harris - talked about this approach. He teaches at Santa Clara University I believe.

2. I received Amy's email about using cheap filters for the high power protect subsystem. I agree that the proposed design was overkill and very expensive. One might consider using a simple diode detector instead of the log detector. Follow the diode detector with an analog amp and a Maxim programmable comparator or use the ADC you planned. Unless I'm missing something you are looking for large signals related to faults and a huge dynamic range is unnecessary. Chris used this approach on PEP-II and it seems quite acceptable. In any case I would recommend not using trim pots in general, it's just another thing requiring adjustment and has potential for error.
3. I think we all agreed that the phase detector/mixer circuit does not belong in the HPPS system, besides that information is available in the RF electronics. By the way I looked at the WJ HMJ5 mixer Tony is using. It does indeed require a +17dBm LO - it's not an active mixer. You might consider using active mixers from RF Micro Devices or (Agilent!) to reduce the amount of high level LO there is running around the board (also gets rid of a substantial LO distribution amplifier). I'm not an expert on these but it might be worth a quick peek. I would encourage Tony to try using striplines between ground layers to transport RF. I think Echotek (sp?) uses this with great success. Beware this requires buried vias but I don't think that's so uncommon these days.
4. OK I looked over the schematics and don't see anything that jumps out at me. I believe you have addressed the issues we talked about last time. What ever you decide to do with the FIR leave some room in the gate array so you can expand the design if something unforeseen comes up. In general I think this modular highly programmable design will allow you to tweak the software and firmware to improve performance as you get some operating experience. One nice upgrade would be to reduce the ramp-up time on each pulse through some serious feed-forward techniques. The NLC proposal is based in this. This would directly improve the wall power efficiency, coming from California with threats of roaming blackouts this now seems more important to me!

Once again thanks for having me out. I really enjoy collaborating with the Los Alamos team, Chris, and Curt. Amy you have done a great job getting the communications channels functioning and creating a positive work environment for you people. I anticipate this project being very successful. Feel free to share this note with your group. Let me know if you have questions.

Tallyho and keep up the pace - the delivery date is fast approaching.

Specific Comments from Chris Ziomek

Overall, the design concept and implementation looks very good. I still am concerned about the schedule, but the reduction in extraneous features from the PDR is a positive step. The LLRF team looks strong and is making good progress. It is too bad that the meeting was cut short, because I have concerns about prototyping, engineering test, manufacturing, etc. (See item 4.)

1. I get the impression that there may be more information that could be achieved through modeling. To-date, the modeling has been used to provide the basic framework for the system and try a number of different topologies, but I think that it should also be used to finalize some of the more important details. For example, how are you planning to dampen the 1 MHz cavity mode? How many taps and what type of filter is really necessary in the feedback loop? What type of compensator bandwidth is necessary in the forward loop? What frequency characteristics are necessary for each filter. What are the FIR/IIR filter coefficients? What delay is actually acceptable for a 200 kHz bandwidth and is this achievable given the desired filter characteristics?
2. Paul and I have discussed the filter characteristics by e-mail and agree (I think) that the FIR filter in the feedback path should be carefully evaluated. The purpose of this filter is to filter out the Fs/2 spurs caused by offset in the I/Q demodulation process. We also agree that the 1 MHz mode should not be filtered from the feedback (as you might want to measure this signal), but filtered from the forward path compensator. Just filtering the mode from the feedback does not guarantee that it will not be excited (by feedforward, etc.) Paul also points out that you might actually want to actively damp this mode with feedback at the proper phase (again something to model). Before Mark spends a lot of time implementing a 20-tap FIR filter, I strongly suggest that you look at a simple 2^{nd} order IIR for this feedback filter which will significantly reduce your processing delay and increase your bandwidth.
3. I worry about the time allotted for DSP firmware. Although flexible, we have found that DSP firmware implementation and debugging can take as long as the hardware implementation. I am not sure that this is accommodated in your schedule.
4. Although not specifically covered in the FDR, I have concerns about the prototyping, design qualification, and manufacturing. For example, are the PCBs expected to be perfect after the first cut? Is there money and time allocated for revisions? Will you have a fully functional system to work out all hardware and firmware issues before having to commit to large-scale manufacturing? Are the designs being implemented as production-ready, fully-documented designs?